LA-UR-87-3675

LUM 8104153--2

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W 7405-ENG-36

TITLE STATUS REPORT ON THE ANSI 100-MBYTE/S CHANNEL

LA-UR--87-3675

DE88 003183

AUTHOR(S) JOHN F. MORRISON

SUBMITTED TO CRAY USERS GROUP MEETING,
BOLOGNA, ITALY
SEPTEMBER 22-26, 1987

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

FORM NO. 836 R4
ST. NO. 2629 5/81

STATUS REPORT ON THE ANSI 100-MBYTE/S CHANNEL

John F. Morrison*
Los Alamos National Laboratory
Los Alamos, New Mexico

INTRODUCTION

Los Alamos National Laboratory has submitted to the American National Standards Institute (ANSI) X3T9.3 committee a specification for a 100-Mbyte/s channel called the high-speed channel (HSC). This channel specification is a simple, yet high-performance, physical layer protocol and media definition. This report will describe the current definition of the HSC, possible applications, and status of our proposal before the ANSI committee.

Higher-speed channels (those at speeds of 100 Mbyte/s) are just beginning to appear on vendor equipment. Cray Research Inc. has a 100-Mbyte/s channel called HSX, Convex Computer Corp. has an 80-Mbyte/s channel called HSP, and other vendors have higher-speed channels under development. In addition, Ultra Corp. is developing a networking system to support channels with these data rates. Each of these vendor-developed higher-speed channels are proprietary in nature. Furthermore, no vendor has proposed a channel standard.

There are many benefits to users of supercomputers for a standard high-speed channel, and it appears that we have a window of opportunity for developing this standard. Consequently, Los Alamos advocated a high-speed channel standard in the ANSI X3T9.3 committee during the past year. At the present time, industry interest is very high. We have a mailing list of 120 people representing 80 companies. In addition, the HSC is being prototyped at IBM, Los Alamos, and other organizations.

GOALS

The primary goal for the HSC is that it should be a fast, digital information delivery system that is optimized for predictable transfers of blocks of data. The HSC design should be simple, provide error detection, support networking applications and easily adapt to fiber-optic implementations. In addition, the HSC design should support data bus sizes of 16, 32, and 64 bits to accommodate 50-Mbyte/s, 100-Mbyte/s, and 200-Mbyte/s implementations.

*This work was performed under the auspices of the U.S. Department of Energy.

FEATURES OF HSC

The HSC proposal currently being considered is a simple physical level description and has the following features: It supports point-to-point connections at 100-Mbyte/s burst transfer rate over a 32-bit-wide data bus. The distance limitation is 25 meters for copper implementations. For fiber implementations, we anticipate supporting distances of tens of kilometers. Error detection is provided by vertical parity (one parity bit for each 8 bits of data) and longitudinal parity (32 bits of check data at the end of a burst). The HSC is full duplex and symmetrical and therefore consists of a transmit channel and a receive channel that are independent of each other. Each channel supports a request information field and has no bidirectional signal lines. These features are important for potential networking switching applications. The HSC is also easily implemented from readily available components. For example, the connecting cables are standard IPI (intelligent peripheral interface) cables.

DESCRIPTION OF HSC

A simple diagram of the HSC is shown in Figure 1. This diagram illustrates the signal lines for the transmit portion. The receive portion is identical. The signal lines consist of a data bus (lines D00 - Dxx), error control bus (P0 - Px), packet control lines, and burst control lines. The data bus can be 16, 32, or 64 parallel lines. The error bus can be 2, 4, or 8 parallel lines. Each error line is associated with 8 data lines and contains the odd parity bit.

The packet control lines control the flow information on the HSC. Information is organized into a hierarchy of packets, bursts, and data elements. A packet is a sequence of bursts, a burst is a sequence of 256 data elements, and a data element is either 16, 32, or 64 bits, depending upon the width of the data bus. The source uses the request signal line to notify the destination that a logical connection is desired. The destination uses the connect signal line in response to a request signal. The source may optionally place a data element on the data bus while asserting the request signal. The meaning associated with this data element is defined by

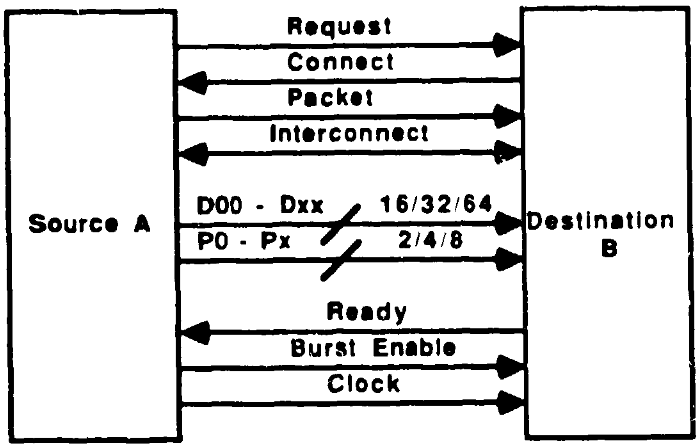

Figure 1. HSC functional diagram.

the next higher protocol layer. One possible use of this is to provide an address for switching applications. Once the source has asserted request and the destination has responded by asserting connect, a logical connection is established. While a logical connection is established, multiple packets may be transmitted from the source to the destination. Each packet is delineated by the packet signal. One additional signal, the interconnect signal, is included in the packet control lines. This signal indicates to the source and destination that the media cable is electrically connected.

The burst control lines consist of ready, burst enable, and clock. The ready signal is used by the destination to notify the source that the destination can accept a burst. The destination can send multiple (pre-acknowledge) ready signals to enable high-transfer rates to be sustained. The burst enable line is used to delineate a burst. The clock is used to synchronize the flow information and currently is defined with a period of 40 ns. Further detail on the definition of the HSC can be obtained from a document prepared by Gene Dornhoff, Los Alamos National Laboratory, entitled *Los Alamos High Speed Channel Proposal*.

APPLICATIONS

The potential applications of the HSC can be discussed at great length. However, the application areas of most interest at Los Alamos are high-speed graphics and networking. At Los Alamos, researchers generate "movie" files from production codes that run several hours on CRAY X-MPs. The researcher views these movie files on a graphics terminal after a production code has finished. Running these movie files at near-video rates (24 frame/s) requires a 50-Mbit/s transfer rate to view a 512- by 512- by 8-bit resolution frame. Many of our users require 1024- by 1024 resolution, which increases the data rate by a factor of four to 200 Mbit/s. A few users would like to see 1024- by 1024- by 24-bit resolution, which increases the data rate to 600 Mbit/s. In addition to viewing results from production runs, researchers also would like to observe the operation of their algorithms in real time. This application was advocated by John von Neumann in his proposal for numerical experiments in 1946!

We have provided a few users at Los Alamos with higher-speed graphics capability. We developed a frame buffer and connected it to a CRAY X-MP/416. The frame buffer operates at 50 Mbit/s. We also have an ultra speed graphics facility that can display images at 60 Mbit/s. The ultra speed graphics facility was described in the Summer 1987 issue of *Cray Channels*. Our challenge is to provide these data rates to a large number of our users through a network.

The Los Alamos Integrated Computer Network (ICN) is illustrated functionally in Figure 2. The various computational servers of the network, the CTSS, NOS, and UNIX hosts, are connected to a variety of network servers such as Common File System (CFS), Print And Graphics Express Station (PAGES), and Facility for Operations Control and Utilization Statistics (FOCUS) by a packet-switched network called the File Transport Network. The XNETs, gateways to other networks, are also connected to the File Transport Network. The File Transport Network accommodates large packet sizes (25 Kbytes) and data rates up to 40 Mbit/s. Access to a large terminal network is provided by another packet switched network.

To support graphics applications and anticipated distributed computing applications, we will require a network that can interconnect a variety of computer systems from different vendors and that can operate at 800-Mbit/s burst data rates. We anticipate networking products at these data rates in the near future. Ultra Corp. is developing networking systems that support these data rates.

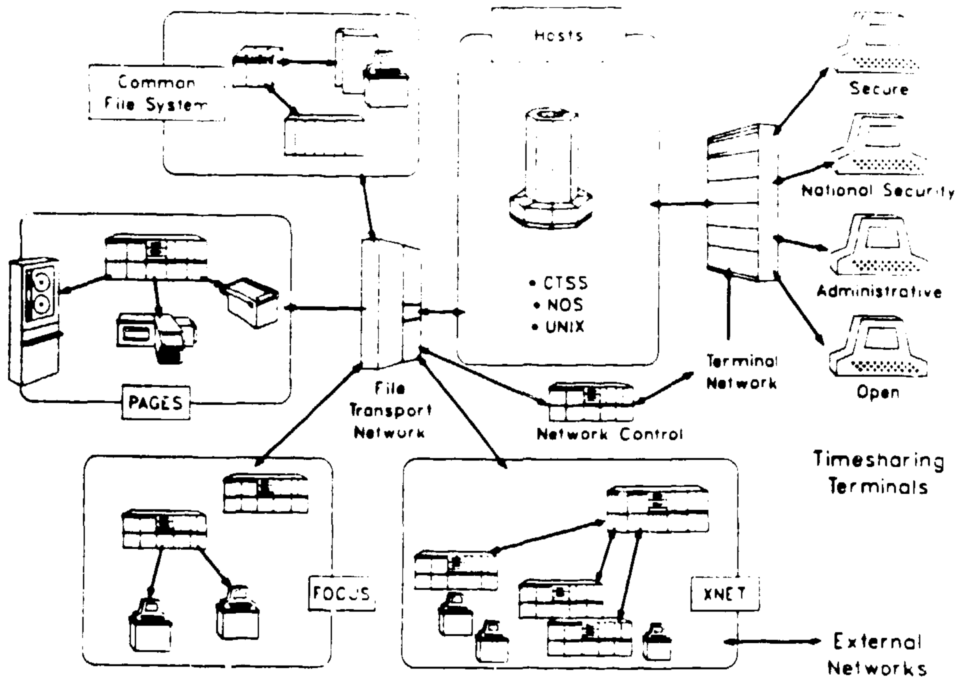

Figure 2. The ICN at Los Alamos.

We are prototyping a crossbar switch that would support these data rates. The development and adoption by industry of a standard high-speed channel would benefit both the users of high-speed networks and those building the networks. Instead of expending resources to develop interfaces to each vendor's proprietary high-speed channel, network vendors could develop a single standard interface and concentrate their resources on networking issues.

CURRENT STATUS

The status of the HSC proposal before the X3T9.3 committee can be briefly summarized. Because the committee is convinced channels of these speeds are needed, two implementations are being considered. The HSC is oriented towards computer-to-computer applications. The other proposal being considered is called IPI Wider-Wider/Faster-Faster. This proposal is an extension of the current IPI standard and is better suited for computer-to-peripheral applications. We believe that both implementations are useful.

Some of the issues that are currently being discussed by the committee are fiber-optic implementations, types of connectors, and the 50-Mbyte/s and 200-Mbyte/s options. The discussions on connectors are addressing concerns about the amount of backplane space required by the connectors. After the requirements for the 50- and 200-Mbyte/s options have been examined, these options may be dropped.

The HSC proposal has attracted interest in industry. In addition to the standards efforts, which could be described as a "top-down" approach, several companies are developing prototypes. This latter approach, which is more "bottom-up," helps to sustain interest in the standards committee.